Remarks

Claims 1-20 are pending in this application. The Examiner rejected claims 11-14 under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent No. 6,081,807 to Story *et al.* (Story). The Examiner rejected claims 1-6 and 16-19 under 35 U.S.C. § 103(a) over Story in view of U.S. Patent No. 6,029,168 to Frey (Frey). The Examiner rejected claims 7-10 and 20 under 35 U.S.C. § 103(a) over Story and Frey in further view of U.S. Patent No. 5,991,763 to Long. The Examiner rejected claim 15 under 35 U.S.C. § 103(a) as being unpatentable over Story in view of Long. Applicants respectfully disagree with the Examiner's rejections.

1. Rejection of Claim 11

The Examiner rejected claim 11 under 35 U.S.C. § 102(e) as being anticipated by Story. Independent claim 11 provides for a method for accessing a file referenced by a file name. The file name is sent to a name server. A file identifier corresponding to the file name is received from the name server. The file identifier is sent to a location server. *The location server is separate from the name server.* File location information corresponding to the file identifier is received from the location server. The file is accessed using the location information.

Story discloses interfacing with a stateless NFS server. As illustrated in Story's Figure 1, network client 102 communicates with network server 106 using network 110. Network client accesses local disk storage 120 and network server 106 accesses network storage 132. As such, Story discloses a system similar to the prior art system illustrated in Applicants' Figure 1 and described by Applicants from page 5, line 8, through page 6, line 2, reproduced as follows (emphasis added):

> Referring to Figure 1, a schematic diagram illustrating a prior art client-server relationship is shown. A file system, shown generally by 20, includes clients 22 accessing data held in files on one or more storage devices 24. Each storage device 24 is accessed through a server, one of which is indicated by 26. A typical example of file system 20 is the Unix-based *Network File System (NFS)*. In NFS, client 22 wishing to access data first *forwards the name of the file containing the data to server 26*, as shown by 28. Server 26 returns a handle to the requested file as indicated by 30. Client 22 then *forwards a data request with*

> *the received handle to server 26*, as indicated by 32. Server 26 requests the data from storage device 24, as indicated by 34. Storage device 24 returns the data to server 26, as indicated by 36, and server 26 forwards the data to client 22, as indicated by 38. Another typical example of file system 20 is embodied in the CIFS (Common Internet File System) found in the WINDOWS NT ® server by Microsoft Corp. Server 26 provides a file descriptor in response to a name provided by client 22. Client 22 uses the file descriptor to access data through a logical connection through server 26 to storage device 24. The file descriptor is valid only for the life of the connection.
>
> There are several problems associated with the traditional client-server system. First, server 26 may not have sufficient resources to support an increasing number of clients 22. Second, the failure of server 26 makes storage device 24 inaccessible by clients 22. Third, a client 22 not directly connected to server 26 may have difficulty locating and accessing a file stored on storage device 24 connected to server 26. Finally, server 26 may not be able to properly respond to client 22 requesting a file using a naming scheme different than the scheme used by server 26.

Story fails to anticipate Applicants' claim 11 for at least the following four reasons.

a. Story doe not disclose separate name server and location server

The Examiner has identified Story's name server 130 as Applicants' name server. The Examiner has also identified Story's NFS server 122 as Applicants' location server. Claim 11 provides that these servers are separate. However, Story discloses name server 130 and NFS server 122 both in network server 106. (*See*, col. 4, ll. 4-14.) Thus, Story neither teaches nor suggests separate name and location servers.

b. Story doe not disclose a name server receiving a file name and sending a file identifier corresponding to the file name

Claim 11 provides for sending a file name to the name server and receiving a file identifier corresponding to the file name from the name server. The Examiner identified Story's name server 130 as Applicants' name server. Story's *entire* disclosure for name server 130 is provided in column 5, lines 7-8, as follows:

> [N]ame server 130 is responsible for file name hierarchy and provides pathname resolution.

Thus, Story neither teaches nor suggests a name server that receives a name and returns a file identifier corresponding to that name.

c. Story doe not disclose a location server receiving a file identifier and sending file location information

Claim 11 provides for sending the file identifier to a location server and receiving file location information corresponding to the file identifier from the location server. The file is then accessed using the file location information. The Examiner identified Story's NFS server 122 as Applicants' location server. Story discloses that NFS server 122 receives a file handle including a file set ID and a file ID in column 4, lines 41-47. Story also discloses that the NFS server forwards the file set ID and file ID to interface 126 at column 5, lines 2-6. Thus, if the file handle is the file identifier, Story's NFS server 122 receives and sends a file identifier and does not send file location information.

d. Story doe not disclose a file identifier received from a name server and sent to a location server

Claim 11 provides for receiving a file identifier corresponding to the file name from the name server and sending this file identifier to the location server. Thus, whatever the Examiner identifies as the file identifier must come from whatever the Examiner identifies as the name server and be sent to whatever the Examiner identifies as the location server. The Examiner identified Story's file ID as Applicants' file identifier. There is no disclosure in Story of any file ID that is received from Story's name server 130 and is sent to Story's NFS server 122.

2. Rejection of Claims 12-15

The Examiner rejected claims 12-14 under 35 U.S.C. § 102(e) as being anticipated by Story. The Examiner rejected claim 15 under 35 U.S.C. § 103(a) as being unpatentable over Story in view of Long. Claims 12-15 are dependent upon claim 11 and, since claim 11 is patentable, are also patentable.

Claim 12 provides that each file is stored as at least one file extent. Story does not disclose storing a file as one or more extents, nor does the Examiner mention any extents in his rejection.

Claim 13 provides that each file is represented in storage as an object and each file identifier is an object identifier. The Examiner asserts claim 13 is disclosed by Story as follows:

> As to Claim 13, Story details a system which including [*sic*] 'each file is represented in storage as an object and each file identifier is an object identifier' [col 4 lin 53-60, col. 5, line 2-6], examiner interpreting object identifier corresponds [*sic*] to Story's VNODE because, OSS file system has the ability to locate VNODE associated with the file using hashing mechanism based on the file ID and the file including file handle as detailed in col 5 line 30-34.

Hashing may be defined as follows:

> **hashing**: A method of transforming a search key into an address for the purpose of storing and retrieving items of data. The method is often designed to minimize the search time.[1]

Thus, hashing has nothing whatsoever to do with objects. Further, Story never mentions or suggests objects in any manner.

3. Rejection of Claim 1

The Examiner rejected claim 1 under 35 U.S.C. § 103(a) over Story in view of Frey. Claim 1, as amended, provides for a file system for storing data. A plurality of storage devices store at least one copy of at least one file. At least one location server maps a file identifier for each file into the location of each copy of the file represented by the file identifier. At least one name server maps a file name to the file identifier referenced by the

---

[1]IBM Dictionary of Computing, 10th Ed., McGraw-Hill, Inc., 1994, pg. 309.

file name. Each name server is physically separate from the at least one location server. No combination of Story and Frey teach or suggest Applicants' claim 1 for at least the following three reasons:

a. ***Story doe not teach or suggest Applicants' name server***

Claim 1 provides for at least one name server operating with file names and file identifiers referenced by the file name. The Examiner identified Story's name server 130 as Applicants' name server. Story's *entire* disclosure for name server 130 is provided in column 5, lines 7-8, as follows:

> [N]ame server 130 is responsible for file name hierarchy and provides pathname resolution.

Thus, Story neither teaches nor suggests a name server that receives a name and returns a file identifier corresponding to that name.

b. ***Story does not teach or suggest Applicants' location server***

Claim 1 provides for at least one location server operating with a file identifier for each file and the location of each copy of the file represented by the file identifier. The Examiner identified Story's NFS server 122 as Applicants' location server. Story discloses that NFS server 122 receives a file handle including a file set ID and a file ID in column 4, lines 41-47. Story also discloses that the NFS server forwards the file set ID and file ID to interface 126 at column 5, lines 2-6. Thus, if the file handle is the file identifier, Story's NFS server 122 receives and sends a file identifier and does not send file location information.

c. ***Neither Story nor Frey teach or suggest separate name server and location server***

Claim 1, as amended, provides that each name server be physically separate from any location server. As seen in Figure 5, whatever mapping Frey discloses is accomplished in data node 42. Thus, Frey does not disclose separate servers for mapping names and file identifiers. Further, the Examiner identified Story's name server 130 as Applicants' name server and identified Story's NFS server 122 as Applicants' location server,

both of which reside in Story's network server 106. (*See*, col. 4, ll. 4-14.) Thus, neither Story nor Frey teach or suggest Applicants' separate name server and location server.

4. Rejection of Claim 16

The Examiner rejected claim 16 under 35 U.S.C. § 103(a) over Story in view of Frey. Claim 16, as amended, provides for a file system for storing data. A plurality of storage devices store at least one copy of at least one file. At least one location database has a map between a file identifier for each file and location information for each copy of the file represented by the file identifier. At least one name database has a map between a file name and the file identifier referenced by the file name. Each name database is physically separate from any location database. At least one client requests a file identifier corresponding to a requested file name. The file identifier mapped to the requested file name is received. Location information corresponding to the received file identifier is requested. Location information mapped to the received file identifier is received. Data is accessed using the location information. No combination of Story and Frey teach or suggest Applicants' claim 16 for at least the following four reasons:

a. Story doe not teach or suggest Applicants' name database

Claim 16 provides for at least one name database operating with file names and file identifiers referenced by the file name. The Examiner identified Story's name server 130 as Applicants' name database. Story's *entire* disclosure for name server 130 is provided in column 5, lines 7-8, as follows:

> [N]ame server 130 is responsible for file name hierarchy and provides pathname resolution.

Thus, Story neither teaches nor suggests a name database that receives a name and returns a file identifier corresponding to that name.

b. Story does not teach or suggest Applicants' location database

Claim 16 provides for at least one location database operating with a file identifier for each file and the location of each copy of the file represented by the file

identifier. The Examiner identified Story's NFS server 122 as Applicants' location database. Story discloses that NFS server 122 receives a file handle including a file set ID and a file ID in column 4, lines 41-47. Story also discloses that the NFS server forwards the file set ID and file ID to interface 126 at column 5, lines 2-6. Thus, if the file handle is the file identifier, Story's NFS server 122 receives and sends a file identifier and does not send file location information.

c. ***Neither Story nor Frey teach or suggest separate name database and location database***

Claim 16, as amended, provides that each name database be physically separate from any location database. As seen in Figure 5, whatever mapping Frey discloses is accomplished in data node 42. Further, the Examiner identified Story's name server 130 as Applicants' name database and identified Story's NFS server 122 as Applicants' location database, both of which reside in Story's network server 106. (*See*, col. 4, ll. 4-14.) Thus, neither Story nor Frey teach or suggest Applicants' physically separate name database and location database.

d. ***Neither Story nor Frey teach or suggest Applicants' client***

Claim 16 provides for a client operative to request a file identifier corresponding to a requested file name, receive the file identifier mapped by a name database to the requested file name, request location information corresponding to the received file identifier, receive location information mapped by a location database to the received file identifier, and access data using the location information. The Examiner asserts that Story's network client 102 is Applicants' client. However, Story discloses receiving data and/or status as a direct result of sending a request. (*See*, Figure 2.) No location information passes to network client 102.

5. Rejection of Claims 2-10 and Claims 17-20

The Examiner rejected claims 2-6 and 17-19 under 35 U.S.C. § 103(a) over Story in view of Frey. The Examiner rejected claims 7-10 and 20 under 35 U.S.C. § 103(a) over Story and Frey in further view of Long. Claims 2-10 depend from claim 1 and, since

claim 1 is patentable, claims 2-10 are also patentable. Claims 17-20 depend from claim 16 and, since claim 16 is patentable, claims 17-20 are also patentable.

Claims 2 and 17 provide that each file is stored as at least one file extent, the file identifier comprising a file handle. The Examiner relies on Story to reject claims 2 and 17. However, Story does not disclose storing a file as one or more extents, nor does the Examiner mention any extents in his rejection.

Claims 3 and 18 provide that each file is represented in storage as an object and each file identifier is an object identifier. The Examiner asserts that such objects are disclosed by Story's hashing. As discussed above, hashing neither teaches nor suggests objects.

Claim 6, as amended, provides that at least one name server operates under a first file access standard and at least one name servers operates under a second file access standard different from the first file access standard. The Examiner indicates NFS as a first standard in Story. The Examiner asserts that a second file access standard is disclosed in Story at column 6, line 64, through column 7, line 12, as follows:

> The file system obtains the information about the current access state of the file from an associated VNODE. The current access state may be read-only, read-and-write, or closed.
>
> At step 504, the file system determines whether the current state of the file matches the needed state indicated by the read or write request. If there is a match, then no further processing is needed. The file already has the appropriate state for the requested operation. If the current file state is not the same as the needed state, the file system then determines whether the current file state is a read and write state at step 508. If so, the file again has the appropriate state for the requested operation and the processing is completed. However, if the current file state is not a read and write state, the file system determines whether the current file state is closed at step 510.

This passage has nothing whatsoever to do with different file access standards. Story neither teaches nor suggests a second standard, let alone the ability for name servers to operate under different file access standards.

Claim 7 provides for at least one client requesting a file identifier for a new file from a location server, receiving the requested file identifier, and registering the file identifier and a new file name for the new file with at least one name server. The Examiner has

provided no teaching or suggestion for a client requesting a file identifier or for such a client registering the file identifier with a new file name.

Conclusion

Claims 1-20, as amended, are pending in this Application. The case is in appropriate condition for allowance. Accordingly, such action is respectfully requested. A check in the amount of $110.00 is enclosed to cover the Petition Fee. Any additional fees may be charged to Deposit Account 19-4545 as specified in the Application Transmittal.

The Examiner is invited to telephone the undersigned to discuss any aspect of this case.

Respectfully submitted,

MARK A. BAKKE et al.

By _____

Mark D. Chuey, Ph.D.
Reg. No. 42,415
Attorney/Agent for Applicant

Date: November 5, 2001

BROOKS & KUSHMAN P.C.
1000 Town Center, 22nd Floor
Southfield, MI 48075
Phone: 248-358-4400
Fax: 248-358-3351

Attachment

VERSION WITH MARKINGS TO SHOW CHANGES MADE

1. (Amended) A file system for storing data comprising:

a plurality of storage devices, each storage device operative to store at least one copy of at least one file;

at least one location server operative to map a file identifier for each file into the location of each copy of the file represented by the file identifier; and

at least one name server operative to map a file name to the file identifier referenced by the file name, each name server physically separate from the at least one location server.

4. (Amended) A file system as in claim 1 wherein each location server [database] is further operative to store metadata associated with each file identifier.

5. (Amended) A file system as in claim 1 wherein at least one location server [database] is on a first computer system and at least one name server [database] is on a second computer system in communication with the first computer system.

6. (Amended) A file system as in claim 1 wherein the at least one name server [database] is a plurality of name servers [databases], at least one of the plurality of name servers [databases] operating under a first file access standard and at least one of the plurality of the name servers [databases] operating under a second file access standard different from the first file access standard.

16. (Amended) A file system for storing data comprising:

a plurality of storage devices, each storage device operative to store at least one copy of at least one file;

at least one location database comprising a map between a file identifier for each file and location information for each copy of the file represented by the file identifier;

at least one name database comprising a map between a file name and the file identifier referenced by the file name, each name database physically separate from the at least one location database; and

at least one client operative to

(a) request a file identifier corresponding to a requested file name,

(b) receive the file identifier mapped to the requested file name,

(c) request location information corresponding to the received file identifier,

(d) receive location information mapped to the received file identifier, and

(e) access data using the location information.